T1C

GESTURE RECOGNITION

APPLICATION NOTE
SEPTEMBER 15, 2025

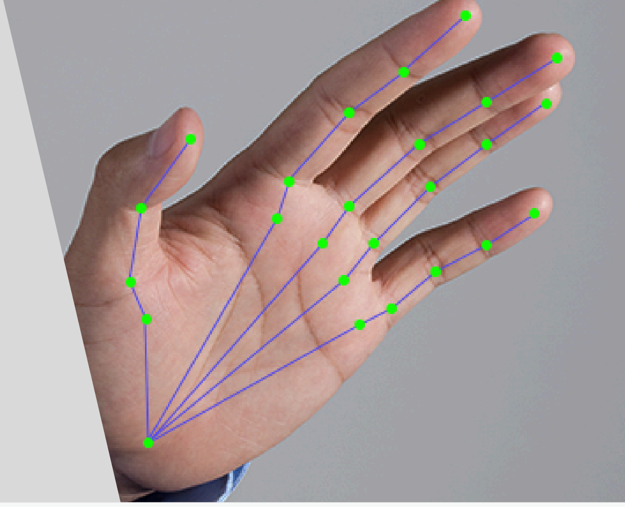

Platform Goal

Gesture recognition for safety-critical Human Machine Interface (HMI) in industrial and clinical settings requires sub-10 ms end-to-end latency, low energy per inference, and robust operation under noise and variable lighting. Real-time response and low power consumption are key requirements that can be effectively addressed by brain-inspired algorithms and neuromorphic systems. To meet these demands, we are developing a spiking neural network (SNN) for classifying gesture recognition data. The project aims to build a real-time gesture recognition model optimized for edge deployment, balancing power efficiency with rapid inference. Evaluation is performed on our T1C custom neuromorphic hardware, which leverages brain-inspired spiking neural networks to process gesture data quickly while minimizing energy consumption. This architecture leverages sparse, recurrent dynamics to preserve temporal micro-structure while minimizing memory bandwidth and power, making it suitable for power-constrained, real-time deployments in life-critical and industrial environments.

Use Case

To evaluate our model, we use the DVS128 Gesture Dataset, which provides event-based recordings of 11 distinct hand gestures performed by 29 subjects across three different lighting conditions, reflecting realistic environmental variability. Each gesture recording spans 6 seconds, with events captured asynchronously by a Dynamic Vision Sensor that detects changes in pixel brightness, offering a high temporal precision with time steps reduced to milliseconds. This rich temporal data allows spiking neural networks to process dynamic, time-dependent information effectively, enhancing their ability to learn and recognize complex gesture patterns. The combination of temporal detail and real-world variability makes this dataset ideal for benchmarking advanced gesture recognition models designed for practical, safety-critical tasks.

Architecture

The spiking recurrent neural network (SRNN) was used because its recurrent architecture captures temporal dependencies essential for dynamic gesture recognition, as shown in Figure 1, where hidden layers have recurrent connections that allow processing of sequences and maintain continuity in spatiotemporal information. This is crucial for gestures with subtle timing differences in event-based DVS data. Although spiking neural networks usually use Leaky Integrate-and-Fire (LIF) neurons, the chosen model replaces them with Liquid Time Constant (LTC) neurons, which dynamically adapt membrane and threshold time constants based on inputs and network states. This enhances the network's ability to model diverse temporal dynamics with fewer layers and less computation, improving training efficiency and making inference suitable for real-time, low-power embedded devices. The network topology includes input neurons corresponding to the frame size, two recurrent layers of 256 neurons each, and classify 11 classes. The total number of timesteps was set to 50.

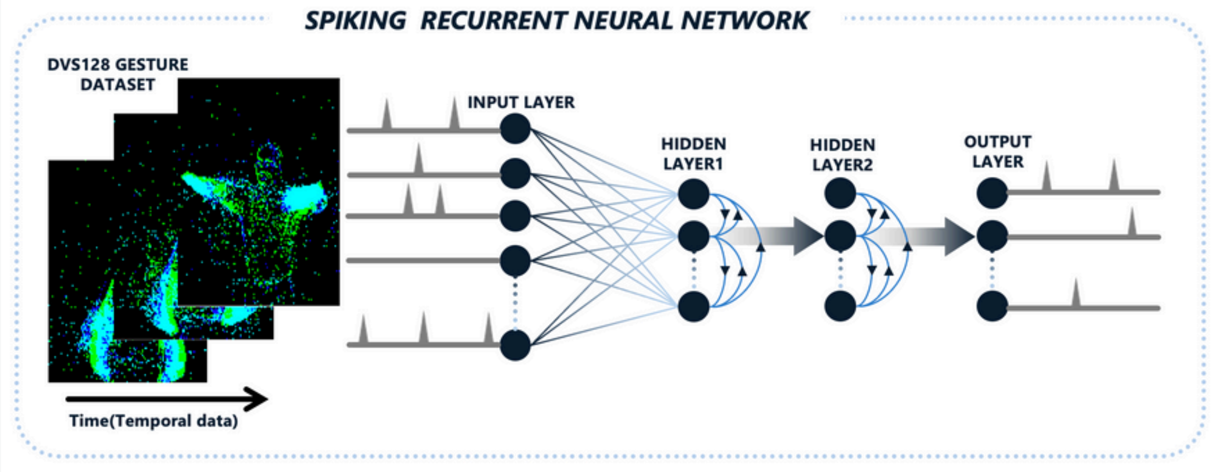

Figure 1: Spiking Recurrent Neural Network Architecture for Gesture recognition

Results	
Device	Power Efficiency GOP/s/W
T1C Neuromorphic Hardware	75.76
NVIDIA Jetson Nano	8.00
NVIDIA RTX 3060 GPU	3.81
Intel i9-12900H CPU	0.31

Technical Performance

Our neuromorphic system, implemented on a Xilinx/AMD Artix UltraScale+ (XCAU15P) FPGA, delivers 41.67 GOP/s at ~0.55 W, achieving 75.76 GOP/s/W. This is **~244× more energy-efficient than an Intel i9-12900H** (0.31 GOP/s/W), **~19.9× more efficient than an NVIDIA RTX 3060** (3.81 GOP/s/W), and **~9.5× more efficient than Jetson Nano** (8.0 GOP/s/W). The combination of sub-watt power and high throughput supports real-time, on-device gesture classification on DVS128 without cloud dependence, aligned with safety-critical HMI requirements. Together with robust performance on DVS128 under variable lighting and noise, the system directly addresses the core requirements for scalable, efficient, and safety-critical gesture detection at the edge.

Roadmap

Current priorities focus on characterizing the baseline gesture recognition model on the custom neuromorphic hardware, despite the current challenges posed by hardware capabilities and configuration within the HLS framework. This initial step will allow us to establish reliable measurements of the model's performance and energy efficiency directly on neuromorphic hardware. Building on these results, the next phase will focus on improving recognition accuracy and energy efficiency by exploring alternative model architectures and training mechanisms, such as more advanced neuron models, optimized learning algorithms, and specialized training strategies tailored for spiking neural networks.

hello@type1compute.com
503.810.8092
www.type1compute.com